

Trading Off Efficiency and Reciprocity in Wireless Peer-To-Peer File Sharing

Fernando Paganini
Universidad ORT Uruguay
Cuareim 1451, Montevideo, Uruguay
Email: paganini@ort.edu.uy

Martín Zubeldía
LIDS, MIT
Cambridge, MA, USA
Email: zubeldia@mit.edu

Andrés Ferragut
Universidad ORT Uruguay
Cuareim 1451, Montevideo, Uruguay
Email: ferragut@ort.edu.uy

Abstract—This paper considers the problem of allocating exchange rates in peer-to-peer dissemination, which must consider the dual objectives of throughput efficiency and reciprocity between peers, the latter essential to cooperation incentives. This question has been studied in prior research for wired networks under an upload constraint, where the focus is on achieving reciprocity through decentralized peer interactions. We consider here a wireless network substrate, for which link capacities are non-uniform according to the peering choice, and exchanges may be subject to interference. A convex optimization problem is formulated that trades off efficiency and reciprocity, and various schemes are investigated to achieve a decentralized solution.

I. INTRODUCTION

Peer-to-peer (P2P) file sharing networks constitute a powerful method for content dissemination in data networks, based on the principle that participating agents are simultaneously clients and servers of a file of common interest, broken up in pieces. While a few servers (“seeders”) are required to sustain the system, the bulk of the capacity can be provided by the client/servers (“leechers”), whose number scales with demand.

The main domain of application of P2P has been the global Internet, where most peers sit behind wired access links. In this context, we highlight two main lines of mathematical research. One concerns the dynamics of peer *populations*, which has been tackled by queueing theory methods [17], [18], [8] as well as fluid differential equation models [13], [12]. These models make assumptions on the service provided by the P2P sharing mechanism, which constitutes the second area of active investigation. Such resource allocation depends on microscopic peer exchange rules, see [7] for a high-level discussion; in particular, a key element are the *reciprocity* mechanisms embedded in P2P protocols to avoid free-riding, such as BitTorrent’s tit-for-tat [5]. More amenable to analytical studies is an alternative *proportional* reciprocity mechanism, studied in [17] and implemented in [11]; the resulting resource allocation is analyzed mathematically in [16], [20].

The above work idealizes the network to capacity constraints in access bandwidth, a suitable assumption for the wired Internet. Far less is known about peer-to-peer performance in other substrates, such as wireless networks. There has been work on epidemic approaches to content dissemination in such networks (e.g. [15]), but here the capacity issue is ignored: exchanges are a one-shot interaction between nodes. In this paper our interest is in a set of nodes who engage in the

dissemination of content over a long period, where the wireless substrate imposes different physical capacities for different links, and links may interfere. There is ample literature on resource allocation in such networks from the network utility maximization perspective (see e.g. [4]); here efficiency and fairness between *sending* flows is considered. This is, however, different from the send/receive *reciprocity* sought here.

The paper is organized as follows. In Section II we set up a general framework for studying efficiency and a convex measure of reciprocity in a P2P network. Optimizing such measure in the wired case was studied in [20]; results are reviewed in Section III. In Section IV we introduce one aspect of the wireless substrate: the fact that outgoing links from each peer will have a non-uniform capacity according to the destination. This introduces an efficiency/reciprocity tradeoff, which is explored by examples, and formalized in terms of a convex optimization problem. Decentralized solutions are devised, which generalize those in the wired case, and are tested in simulation examples. In Section V we tackle the additional issue of link interference. A formalism is introduced to cast our efficiency/reciprocity optimization in this setting, highlighting general properties of the solution. Decentralization is far more challenging here; for a CSMA-type substrate we pursue our tradeoff through a stochastic optimization method similar to [9]. While efficiency is naturally included, reciprocity measures are harder to adapt; we outline a proposal in this direction, deferring details to the thesis [21]. Conclusions are given in Section VI.

II. EFFICIENCY AND RECIPROCITY IN P2P SHARING

We consider a fixed population of N peers who engage in bilateral exchange of information¹. Their connectivity is specified by an adjacency matrix $A = (a_{ij})$, where $a_{ii} = 0$ and $a_{ij} = 1$ when i can send data to j . A is assumed *symmetric*, and with no rows of zeros (no disconnected peers).

Define a *resource sharing* matrix $Z \in \mathbb{R}_+^{N \times N}$ in which z_{ij} represents the offered throughput from peer i to peer j . Z is constrained by connectivity, i.e. it must satisfy

$$z_{ij} \geq 0, \quad z_{ij} = 0 \text{ if } a_{ij} = 0. \quad (1)$$

Z will also be subject to bandwidth constraints, of a nature depending on the network substrate: postponing this question, we refer to a set \mathcal{Z} of feasible allocation matrices.

¹We focus here only on the *leechers*, who both send and receive data.

The aggregate sent and received rates per peer are

$$s_i(Z) = \sum_j z_{ij} \quad \forall i; \quad r_j(Z) = \sum_i z_{ij} \quad \forall j. \quad (2)$$

In matrix form, we can write: $Z\mathbf{1} = s$, $\mathbf{1}^T Z = r^T$, where $s, r, \mathbf{1} \in \mathbb{R}^N$ are interpreted as column vectors, the latter being the vector of ones, and T denotes transpose.

We now proceed to specify the desirable objectives on the resource sharing matrix $Z \in \mathcal{Z}$. A first natural objective is the total rate of the exchange, obtained by

$$R(Z) = \sum_{i,j} z_{ij} = \sum_i s_i = \sum_j r_j. \quad (3)$$

Maximizing this quantity would lead to the most *efficient* use of the network. This objective is however insufficient for peer-to-peer networks; it may happen that the maximum rate allocation leaves a peer receiving no data, which is not compatible with a bidirectional exchange. Instead, to provide proper incentives for cooperation between peers it is desirable to have approximate parity between the rates a peer sends and receives. Such is the fairness criterion chosen by [7], where tradeoffs with efficiency were shown.

We formalize this tradeoff further by using a quantitative measure of discrepancy between the upload and download rate vectors, the Kullback-Leibler (KL) divergence (see [2], [6])

$$D(s||r) := \sum_j s_j \log \left(\frac{s_j}{r_j} \right), \quad (4)$$

a jointly convex function of both vectors. Since s and r have the same sum ($R(Z)$ in (3)), their KL divergence is always non-negative [6] and only zero if $s = r$, i.e. if every peer receives as much throughput as it provides to the network.

Therefore, an allocation matrix Z with small $D(s||r)$ achieves an approximate level of *reciprocity* between a peer and the rest of the swarm. An alternative, *peerwise* notion of reciprocity is to require that peers i and j share equal amounts of bandwidth, i.e. that Z is symmetric. Approximate peerwise reciprocity can be measured by the KL divergence between *matrices*

$$D(Z||Z^T) := \sum_{i,j} z_{ij} \log \left(\frac{z_{ij}}{z_{ji}} \right). \quad (5)$$

The following Lemma formalizes the intuitive fact that the latter notion of reciprocity is more restrictive than the first:

Lemma 1. $D(Z||Z^T) \geq D(s||r)$ for any allocation Z , and equality holds if and only if $\frac{z_{ij}}{z_{ji}} = \frac{s_i}{r_i}$ for any i, j with $a_{ij} = 1$.

Proof. The result follows from the log-sum inequality (see [6]), which for each fixed i implies

$$\sum_j z_{ij} \log \left(\frac{z_{ij}}{z_{ji}} \right) \geq \left(\sum_j z_{ij} \right) \log \left(\frac{\sum_j z_{ij}}{\sum_j z_{ji}} \right) = s_i \log \left(\frac{s_i}{r_i} \right).$$

Adding over i gives the desired bound. Conditions for equality also follow from those in [6]. \square

The question of interest is, under the physical constraints of each specific network scenario, what is a suitable tradeoff between efficiency and reciprocity, and whether such allocation can be found through decentralized peer interactions.

III. WIRED NETWORKS WITH UPLOAD CONSTRAINT

In this section we provide background on prior results for this question in the case of a wired P2P network, under the usual assumption that the only bottleneck is the overall upload bandwidth μ_i from each peer i . In this case we can characterize the allowable resource sharing matrices as

$$\mathcal{Z} = \left\{ Z \in \mathbb{R}_+^{N \times N} \text{ satisfying (1), } \sum_j z_{ij} = \mu_i \quad \forall i \right\}. \quad (6)$$

Here the vector s of total sending rates is fixed at $\mu = (\mu_i)$, and the overall transfer rate is $R(Z) = \sum_i \mu_i$ for all $Z \in \mathcal{Z}$: all allowable allocations are equally efficient.²

Therefore in this case there is no tradeoff, the remaining objective of reciprocity can stated in terms of the following convex optimization:

Problem 1. Given A and μ , find $Z \in \mathcal{Z}$ defined by (6) that minimizes $D(\mu||r(Z))$.

We note again that if $r = \mu$ is feasible within \mathcal{Z} , it will be optimal; otherwise we are seeking a certain kind of approximation. An equivalent formulation (since μ is fixed) is

$$\max_Z \sum_j \mu_j \log (r_j(Z)), \text{ subject to } Z \in \mathcal{Z}.$$

In this version it can be interpreted as an instance of (weighted) *proportional fairness*, extensively studied in Internet resource allocation [10]. Here, we choose each node's weight as its own contribution to the network.

In our recent work [20] we have characterized the set of solutions to Problem 1 through Lagrangian duality, based on a stream of related literature [14], [16]. All solutions Z^* to Problem 1 correspond to a unique vector $r^* = r(Z^*)$, characterized by a unique set of multipliers or prices $p_i^* > 0$, $i = 1, \dots, N$, such that:

- $r_i^* = p_i^* \mu_i$ for every peer. So p_i^* defines the proportional reciprocity the peer receives from the network.
- $z_{ji}^* > 0$ only for $j \in \arg \min \{p_j^* : a_{ji} = 1\}$; furthermore, in this case $p_i^* = [p_j^*]^{-1}$. So at optimality a peer can only receive/send rate to another of *inverse* price.

Also in [20] is a detailed study of a prominent decentralized algorithm for reciprocity, proposed in [17], [16], [11]:

$$z_{ij}(t+1) = \mu_i \frac{z_{ji}(t)}{r_i(t)}. \quad (7)$$

In this *proportional reciprocity* scheme, peer i allocates to peer j the fraction of its bandwidth μ_i equal to the proportion of bandwidth received from peer j in the previous step. In matrix form we write $Z(t+1) = \mathcal{R}[Z(t)]$ by introducing the *reciprocity mapping* $\mathcal{R}[Z] := \text{diag}(\mu_i/r_i(Z)) \cdot Z^T$.

²One could, instead, define \mathcal{Z} by an inequality constraint, but this deliberate inefficiency would serve no purpose, and will not be pursued.

Algorithm (7) is closely related to the so-called Sinkhorn procedure for matrix row and column renormalization [14]. A summary of its main properties is:

- Any solution Z^* of Problem 1 is a fixed point of \mathcal{R}^2 , square of the reciprocity mapping. Furthermore $Z^+ := \mathcal{R}[Z^*]$ is also a solution of Problem 1.
- In general, Z^+ need not be equal to Z^* , i.e. Z^* need not be a fixed point of the map \mathcal{R} itself. However the point $\tilde{Z} = \frac{Z^* + Z^+}{2}$ is another optimum and a fixed point of \mathcal{R} .
- In the special case where $r = \mu$ is feasible, there is always a symmetric optimal allocation.

The most important fact is the following convergence result.

Theorem 2 ([16], [20]). *Given an initial condition $Z(0) \in \mathcal{Z}$ with $z_{ij}(0) > 0$ whenever $a_{ij} = 1$, the sequence generated by (7) satisfies $\lim_{k \rightarrow \infty} Z(2k) = Z^*$, $\lim_{k \rightarrow \infty} Z(2k+1) = Z^+$, where both Z^* and Z^+ are optimal points of Problem 1. Furthermore, $r(Z(t))$ converges to the optimal rate vector r^* .*

Thus, provided initially all exchange options are explored, the even and odd subsequences converge to (possibly different) optimal allocations, and the fairness objective is achieved.

We finish the section by highlighting an additional fact: for any fixed point \tilde{Z} of \mathcal{R} , we have

$$\tilde{z}_{ij} = \frac{\mu_i}{\tilde{r}_i} \tilde{z}_{ji} \implies \frac{\tilde{z}_{ij}}{\tilde{z}_{ji}} = \frac{\mu_i}{\tilde{r}_i} \quad \forall j,$$

the condition for equality in Lemma 1. We conclude that $D(\tilde{Z}||\tilde{Z}^T) = D(\mu||\tilde{r})$. Since as mentioned before there is always a fixed point of \mathcal{R} among the optima of Problem 1, we have the following consequence:

Corollary 3. *The minimum of $D(Z||Z^T)$ under $Z \in \mathcal{Z}$ defined by (6) has the same value as Problem 1, and a subset of its solutions.*

In other words, even if our fairness objective only concerns the *global* reciprocity each peer receives from the network, in this wired network case it is equivalent to optimizing a measure of *peerwise* reciprocity.

IV. WIRELESS NETWORKS: MULTI-RATE PHYSICAL LAYERS

We now move to consider a wireless network substrate, in which peers occupy certain spatial locations, connected by wireless channels. There are at least two differences between this situation and the wired case:

- 1) Wireless channels often adapt their rate to physical layer parameters such as signal-to-noise ratio, affected by distance. As a result, the sending rate will no longer be agnostic to the choice of receiving peer.
- 2) Wireless links may interfere with each other.

In this section we focus on the first issue, postponing the second. So for now we assume all peers have separate transmission channels, which they can allocate independently. Given two peers i and j , let μ_{ij} denote the maximum rate at which peer i can transmit to j , if it chose *only* this destination.

By time-sharing between destinations the peer can achieve the sending rates $z_{ij} = \pi_{ij}\mu_{ij}$ where $\sum_j \pi_{ij} = 1$. Here π_{ij} is the proportion of time devoted by peer i to neighbor j , again we assume no inefficient idle time. This leads to the following set of achievable rate allocations:

$$\mathcal{Z} = \left\{ Z \in \mathbb{R}_+^{N \times N} \text{ satisfying (1), } \sum_j \frac{z_{ij}}{\mu_{ij}} = 1 \quad \forall i \right\}. \quad (8)$$

Note that this is a generalization of (6), which corresponds to the special case $\mu_{ij} = \mu_i$ for all j , where the channel from peer i has the same quality for all destinations.

In general the matrix $M = (\mu_{ij})$ need not be symmetric: differences in peer channel qualities (e.g. transmission power) may cause $\mu_{ij} \neq \mu_{ji}$; indeed asymmetry was already present in the wired scenario.

We now look at a motivating example.

Example 1. *Consider a wireless network with 3 peers which are all neighbors, and the matrix of maximum rates*

$$M = \begin{bmatrix} 0 & 2 & 2 \\ 2 & 0 & 1 \\ 2 & 1 & 0 \end{bmatrix}.$$

Here maximum rates are symmetric, but not uniform among outgoing links. We specify the time-sharing matrix Π and the resulting rate allocation Z :

$$\Pi = \begin{bmatrix} 0 & p & 1-p \\ q & 0 & 1-q \\ v & 1-v & 0 \end{bmatrix}, \quad Z = \begin{bmatrix} 0 & 2p & 2-2p \\ 2q & 0 & 1-q \\ 2v & 1-v & 0 \end{bmatrix}. \quad (9)$$

The set \mathcal{Z} of allowable allocations corresponds to all above matrices Z where p, q, v vary in the interval $[0, 1]$.

The total rate is $\sum_{i,j} z_{ij} = 4 + q + v$, so efficiency is no longer agnostic to the peering choice: the set of efficient allocations is

$$\mathcal{Z}^{\text{eff}} = \left\{ Z = \begin{bmatrix} 0 & 2p & 2-2p \\ 2 & 0 & 0 \\ 2 & 0 & 0 \end{bmatrix}, p \in [0, 1] \right\}.$$

Also note that there are no symmetric matrices in \mathcal{Z}^{eff} , even though M is symmetric.

We now look at reciprocity, computing the vectors

$$s(Z) = \begin{bmatrix} 2 \\ 1+q \\ 1+v \end{bmatrix}, \quad r(Z) = \begin{bmatrix} 2(q+v) \\ 2p+1-v \\ 3-2p-q \end{bmatrix}.$$

It is easily checked that in this case $s = r$ is feasible (thus minimizing $D(s||r)$), achieved for $p = \frac{1}{2}$ and $q + v = 1$. Therefore global reciprocity is reached by the allocations in

$$\mathcal{Z}^{\text{rec}} = \left\{ Z = \begin{bmatrix} 0 & 1 & 1 \\ 2q & 0 & 1-q \\ 2(1-q) & q & 0 \end{bmatrix}, q \in [0, 1] \right\}.$$

Among them, one matrix (for $q = \frac{1}{2}$) is symmetric, thus achieving peerwise reciprocity.

The main observation is that $\mathcal{Z}^{\text{eff}} \cap \mathcal{Z}^{\text{rec}} = \emptyset$, one cannot satisfy both objectives simultaneously.

The example shows that there is a tradeoff between efficiency and reciprocity in wireless P2P settings. This suggests managing the tradeoff through a combined cost that contemplates both factors, such as

$$J(Z) = D(s||r) - \alpha R(Z) = \sum_i s_i \left[\log \left(\frac{s_i}{r_i} \right) - \alpha \right]. \quad (10)$$

Here the parameter $\alpha > 0$ weighs the importance assigned to efficiency. Note that $J(Z)$ is a *convex* function of Z , so its minimization over \mathcal{Z} is a convex optimization problem.

Problem 2. *Given A and M , find $Z \in \mathcal{Z}$ defined by (8) that minimizes $J(Z)$.*

Example 2 (Continuation of Example 1). *We minimize the cost $J(Z)$ over matrices $Z(p, q, v)$ as in (9). We argue that it suffices to confine our search to $p = \frac{1}{2}$, $q = v$. This is because for any point (p, q, v) we can find another point $(p', q', v') = (1-p, v, q)$ with the same efficiency and reciprocity: $R(Z') = R(Z)$, and $D(s'||r') = D(s||r)$, in fact s', r' coincide with s, r modulo a permutation of the last two components. Therefore $J(Z') = J(Z)$. Invoking convexity of J , it can be no larger at the point $(Z + Z')/2$, which corresponds to $p = \frac{1}{2}$, $q = v$.*

We thus consider the scalar valued function in $q \in [0, 1]$:

$$\begin{aligned} \varphi(q) &= J(Z(1/2, q, q)) \\ &= 2 \log \left(\frac{1}{2q} \right) + 2(1+q) \log \left(\frac{1+q}{2-q} \right) - \alpha(4+2q). \end{aligned}$$

Minimizing $\varphi(q)$ does not yield a closed form solution, but we find that the optimal q^ satisfies*

$$\begin{cases} \frac{1}{2} < q^* < 1 & \text{if } 0 < \alpha < 2 + \log(2); \\ q^* = 1 & \text{if } \alpha \geq 2 + \log(2). \end{cases}$$

Thus if the weight α is large we just get the optimal efficiency solution. For moderate values of α we have a compromise with reciprocity which, however, always yields

$$Z(q^*) = \begin{bmatrix} 0 & 1 & 1 \\ 2q^* & 0 & 1 - q^* \\ 2q^* & 1 - q^* & 0 \end{bmatrix}$$

which is non-symmetric and with $s^ \neq r^*$ (since $q^* > \frac{1}{2}$).*

Returning to the general case, we may attempt to solve Problem 2 through duality, which was a powerful method in the wired network situation. Writing the Lagrangian

$$L(Z, \lambda) = J(Z) + \sum_i \lambda_i \left(\sum_j \frac{z_{ij}}{\mu_{ij}} - 1 \right) \quad (11)$$

leads after some analysis to the saddle point condition:

$$\lambda_i^* = \max_{\{j: a_{ij}=1\}} \left[-\log \left(\frac{s_i^*}{r_i^*} \right) + \alpha - 1 + \frac{s_j^*}{r_j^*} \right] \mu_{ij}.$$

In comparison to the conditions reviewed in Section III for the wired case ($\mu_{ij} = \mu_i$) we do not have here a clean interpretation for the optimal multipliers as reciprocity factors. And the preceding coupled transcendental equation does not suggest an immediate path for decentralization.

This motivates us to consider an alternative convex optimization problem:

Problem 3. *Given A and M , find $Z \in \mathcal{Z}$ defined by (8) that minimizes the “energy”*

$$\begin{aligned} E(Z) &= D(Z||Z^T) - \alpha R(Z) \\ &= \sum_{i,j} z_{ij} \left[\log \left(\frac{z_{ij}}{z_{ji}} \right) - \alpha \right]. \end{aligned} \quad (12)$$

In the wired case of Section III, it follows from Corollary 3 that the minimum of $E(Z)$ coincides with that of $J(Z)$ in (10) (in that case the throughput term is constant). In the wireless situation this is no longer true. Still, this alternative of trading off *peerwise* reciprocity with efficiency is a valid option to achieve our tradeoff in a decentralized way, and even leads to new interpretations for the wired case, as studied below.

A. Best response optimization of local energy

To pursue a decentralized solution we first write the energy cost of (12) as

$$E(Z) = \sum_i \overbrace{\sum_j z_{ij} \left[\log \left(\frac{z_{ij}}{z_{ji}} \right) - \alpha \right]}^{E_i(Z_i, Z_{-i})}.$$

Here we are denoting by Z_i the i -th row of Z (i.e. the upload allocations of peer i), and we using a game-theoretic notation in which Z_{-i} denotes allocations of all other peers. An idea for pursuing a decentralized minimization is for each peer to compute the “best response” strategy that minimizes its portion of the cost, given the rest as fixed. This is now investigated.

Problem 4. *For fixed i , and given z_{ji} for all j , minimize $E_i(Z_i, Z_{-i})$ over $Z_i = (z_{ij})$, subject to*

$$\sum_j \frac{z_{ij}}{\mu_{ij}} = 1. \quad (13)$$

To solve this problem we invoke once more a Lagrangian with one multiplier λ_i for the constraint:

$$L_i(Z_i, Z_{-i}, \lambda_i) = \sum_j z_{ij} \left[\log \left(\frac{z_{ij}}{z_{ji}} \right) - \alpha + \frac{\lambda_i}{\mu_{ij}} \right] - \lambda_i.$$

To minimize over Z_i for fixed λ_i (and Z_{-i}), we impose

$$\frac{\partial L_i}{\partial z_{ij}} = \log \left(\frac{z_{ij}}{z_{ji}} \right) - \alpha + \frac{\lambda_i}{\mu_{ij}} + 1 = 0,$$

whose solution gives the reciprocity rule

$$z_{ij} = z_{ji} e^{\alpha - 1 - \frac{\lambda_i}{\mu_{ij}}}. \quad (14)$$

The value of λ_i can be found by imposing the constraint (13).

Remark 1. *An important observation is that in the wired case ($\mu_{ij} = \mu_i \forall j$) we obtain in (14) $z_{ij} = \kappa_i z_{ji}$ for all j , namely a proportional allocation of upload rates as a function of rates received. After imposing the constraint we find $\kappa_i = \frac{\mu_i}{r_i}$ and therefore this solution is precisely the proportional reciprocity iteration of (7). We have thus re-interpreted this algorithm as a best response iteration for the energy cost in (12).*

Motivated by its good properties in the wired case, we test this best response generalization. There are two variants:

- A “Jacobi”-type iteration where all Z_i are updated simultaneously following (14). This is indeed what was done in (7) for the wired case.
- A “Gauss-Seidel”-type algorithm where rows Z_i are updated one-at-a-time.

Example 3. We explore the properties of our algorithms using Matlab simulations, for a 3-node network with maximal rates $\mu_{12} = \mu_{21} = 3; \mu_{13} = \mu_{31} = 2; \mu_{23} = \mu_{32} = 1$.

Figure 1 shows the trajectories of the energy cost $E(Z)$ in (12) when $Z(t)$ is updated using the best-response approach in both variants, Jacobi and Gauss-Seidel; the third “Global Gauss Seidel” algorithm will be described later. We see that the best response algorithms³ are unable to reach the optimum of Problem 3, which is not surprising. Indeed, even in the wired case the corresponding Jacobi iteration (7) does not necessarily converge to optimality in $D(Z||Z^T)$, it can oscillate between two suboptimal points.




Figure 1. Evolution of $E(Z)$ for different algorithms.

Recall that (7) did have the property of reaching the optimal global reciprocity $D(s||r)$. This suggests looking here at its counterpart, the cost $J(Z)$ in (10). However, for this metric the best-response iteration does not perform well: Figure 2 shows that neither version is able to minimize $J(Z)$ as desired.

B. Gauss-Seidel optimization of global energy

In view of the preceding limitations of the best response method, for the wireless case another approach is required. We propose here an alternative that better exploits the underlying potential energy: a Gauss-Seidel algorithm for the global cost $E(Z)$, minimized one row Z_i at a time. Differently from the previous case, the new algorithm considers the effect of Z_i not only in $E_i(Z_i, Z_{-i})$ but also its influence on $E_j, j \neq i$.

This is done as follows. For each fixed i , write the cost in Problem 3 as $E(Z) = \tilde{E}_i(Z_i, Z_{-i}) + \tilde{E}_{-i}(Z_{-i})$, where all

³Often in game theory the term best-response dynamics refers to a local motion along the gradient $\partial_{Z_i} E_i(Z_i, Z_{-i})$; these may converge in certain potential games. Here, however, we are considering a one-shot optimization over Z_i , which has much faster dynamics but weaker guarantees.




Figure 2. Evolution of $J(Z)$ for different algorithms.

terms involving Z_i are included in \tilde{E}_i , and the function \tilde{E}_{-i} depends only on the allocations of other peers. Specifically,

$$\begin{aligned} \tilde{E}_i(Z_i, Z_{-i}) &= \sum_{j:a_{ij}=1} \left[z_{ij} \log \left(\frac{z_{ij}}{z_{ji}} \right) + z_{ji} \log \left(\frac{z_{ji}}{z_{ij}} \right) - \alpha z_{ij} \right] \\ &= D(Z_i || Z_i^T) + D(Z_i^T || Z_i) - \alpha \mathbf{1}^T Z_i, \end{aligned} \quad (15)$$

where Z_i^T is the vector $(z_{ji})_{j \neq i}$, transpose of the i -th column of Z^T . One step in the global Gauss Seidel algorithm is given by:

Problem 5. For fixed i , and given Z_{-i}^T minimize $\tilde{E}_i(Z_i, Z_{-i})$ over $Z_i = (z_{ij})_{\{j:a_{ij}=1\}}$, subject to (13).

This step requires the same information as the best-response version, namely the rates z_{ji} received from other peers, which is the basic assumption of any reciprocity scheme. Both can be computed numerically using convex optimization techniques.

Note that given Z_{-i} , the function $\tilde{E}_i(Z_i, Z_{-i})$ is strictly convex in Z_i , thus Problem 5 has a unique solution Z_i .

Clearly, the Gauss-Seidel iteration will compute a sequence $Z(t)$ with monotonically decreasing values of $E(Z(t))$; will it reach optimality? In the remaining trajectory of Figure 1 we show a simulation for Example 3, which indeed exhibits convergence to the optimum; this behavior is robust to initial conditions. We believe this is a quite general situation:

Conjecture 4. Given an initial condition $Z(0) \in \mathcal{Z}$ with $z_{ij}(0) > 0$ whenever $a_{ij} = 1$, let $Z(t)$ be generated by successive steps of Problem 5, one row at a time. Then any limit point Z^* of $Z(t)$ is a global optimum of Problem 3.

In terms of proving this conjecture, we are “almost” under the conditions of a standard convergence result in [1, Section 3.3.5]. This reference considers a Gauss-Seidel optimization of a smooth convex cost $F(x_1, x_2, \dots, x_n)$ successively in each $x_i \in X_i$ (the global domain is a Cartesian product), and assumes strict convexity on any single variable when the others are fixed. These conditions hold here if one takes each variable to be the row Z_i ; in particular the domain is a Cartesian product since constraints (13) are decoupled in i .

The only technical complication is that the above requires smoothness in the entire domain, and here we have potential pathologies at the boundaries $z_{ij} = 0$, $z_{ji} = 0$, where our cost is not well-defined or tends to infinity. Requiring that $Z(0)$ uses initially all available peering options, the initial problem is well defined and it is not hard to see that this property will be preserved, $z_{ij}(t) > 0$ whenever $a_{ij} = 1$: Problem 5 will avoid the infinite cost (15) at the boundary. However in the limit the boundary *could* be approached, yielding $z_{ij}^* = 0$ (and necessarily, $z_{ji}^* = 0$) for a pair of neighbors; indeed this may be the optimal allocation, avoiding the use of a very inefficient link. Dealing with this case goes beyond the confines of [1], hence we only state a conjecture. We illustrate this situation with another simulation Example.

Example 4. Consider a line network of 4 nodes, with maximal rate and resource sharing matrices M , Z below ($p, q \in [0, 1]$):

$$M = \begin{bmatrix} 0 & 2 & 0 & 0 \\ 2 & 0 & 1 & 0 \\ 0 & 1 & 0 & 2 \\ 0 & 0 & 2 & 0 \end{bmatrix}; \quad Z = \begin{bmatrix} 0 & 2 & 0 & 0 \\ 2p & 0 & 1-p & 0 \\ 0 & 1-q & 0 & 2q \\ 0 & 0 & 2 & 0 \end{bmatrix},$$

A simple analysis reveals that for any $\alpha > 0$ the optimal energy is achieved at $p = q = 1$, i.e. it is optimal not to use the links between peers 2 and 3. Starting our Gauss-Seidel algorithm from a random initial condition that uses all links, we observe convergence. Indeed after 5 rounds of updates the matrix is already using values of p and q of the order of $1 - 10^{-4}$, indicating the undesirable link is being turned off.

While we have yet to acquire extensive experience with this algorithm, in our trials so far we have observed very fast convergence to optimality in terms of number of iterations. Although each iteration involves a convex program, the overall performance is quite promising.

Remark 2. We are not claiming that the alternative cost $J(Z)$ could be also minimized by this algorithm, indeed Figure 2 shows a small but nonzero gap, reflecting a difference between $D(Z||Z^T)$ and $D(s||r)$ at the convergence point. In the wireless case, the equality conditions in Lemma 1 may not hold at the optimum of Problem 3.

V. INTERFERENCE AND MEDIUM ACCESS CONTROL

We now consider a second aspect characteristic of wireless local area networks: the use of a shared medium gives rise to *interference*, preempting certain links from being activated simultaneously. A typical case is when a common wireless channel is used in the network, so nodes interfere with all others within their range. Thus, transmissions must be regulated to avoid interference either by centralized scheduling or by decentralized medium access control.

Interference can generally be characterized through the notion of an *independent set*, a subset of links which are allowed to transmit simultaneously. Consistently with our method of analysis, we will represent sets of active links through an $N \times N$ matrix X , such that $x_{ij} = 1$ if link i to j is active, $x_{ij} = 0$ otherwise. In particular we impose the hard zeros of

the neighborhood structure, $x_{ij} = 0$ whenever $a_{ij} = 0$. We denote by $\mathcal{X} \subset \{0, 1\}^{N \times N}$ the set of matrices that correspond to all independent sets of links.

At any time instant, only one of such configurations can be active⁴. As in the previous section we introduce time-sharing among configurations to achieve a richer mix of transmissions. Let $\pi(X)$ denote the probability (fraction of time) that configuration X is active, with $\sum_{X \in \mathcal{X}} \pi(X) = 1$. Then

$$\Pi = \sum_{X \in \mathcal{X}} \pi(X) \cdot X$$

is the matrix of fractions of time π_{ij} each link is activated. Let \mathcal{P} denote the set of all possible such Π matrices, for different choices of time share distributions $\pi(X)$. Finally, for each link we introduce the maximum transmission rate μ_{ij} ; the effective transmission rate at the link is $z_{ij} = \pi_{ij} \mu_{ij}$ as before. In matrix form $Z = M \circ \Pi$ (componentwise, Hadamard product); again \mathcal{Z} is the resulting set of possible allocation matrices. We note the following general properties.

- \mathcal{P} is the convex hull of \mathcal{X} , and is thus convex.
- \mathcal{Z} is also convex.

Remark 3. We can recast the situation of Section IV in this more general context. In that case the only interfering links are those outgoing from the same peer i , who can only talk to one other peer at once. So our set \mathcal{X} is made of matrices with a single “1” per row, of the structure defined by A . The corresponding set \mathcal{P} are the row-stochastic matrices of structure A , and the set \mathcal{Z} coincides with the one in (8).

Our objective is, as before, to study tradeoffs between efficiency and reciprocity within the set \mathcal{Z} of allowable file-sharing matrices. A natural proposal is to minimize a convex function such as $J(Z)$ in (10) or $E(Z)$ in (12) over \mathcal{Z} , which are convex optimization problems. The challenge, as always, is to achieve this without centralized computation.

A. Symmetry in allocations

The class \mathcal{X} of independent set matrices is said to be *symmetric* if $X \in \mathcal{X} \implies X^T \in \mathcal{X}$; i.e. reversing all peer transmissions does not introduce interference. Note that:

- This does not mean the matrices X themselves are symmetric; links (i, j) and (j, i) will in many situations not be active at once.
- The structure of Section IV is *not* symmetric. For instance in Example 1 we can have links 1-3 and 2-3 active at once, but not the other way round.
- If \mathcal{X} is symmetric (closed under transposition), then so is its convex hull \mathcal{P} .

An important case of a symmetric class \mathcal{X} occurs in the 802.11 (WiFi) standard when the request-to-send/clear-to-send option is activated. Here links (i, j) and (j, i) cannot be on at once. Now, activating a link (i, j) requires a free medium (measured by carrier-sense), and a bidirectional (RTS/CTS) handshake

⁴It is natural to allow inefficient configurations (non-maximal independent sets) due to the difficulty of orchestrating such transmissions among nodes.

between nodes, which establishes that both the forward and backward links are free of interference. In that case, if $x_{ij} = 1, x_{ji} = 0$ is allowed, so is $x_{ij} = 0, x_{ji} = 1$, with the rest unchanged. This is a stronger condition than symmetry of \mathcal{X} .

The following result concerns optimal allocations for symmetric interference sets under the (restrictive) condition of symmetric maximal rates.

Proposition 5. *If \mathcal{X} is symmetric, and $M = (\mu_{ij}) = M^T$, then \mathcal{Z} is a symmetric set, and the optimum of $E(Z)$ in (12) is achieved at a symmetric matrix ($Z = Z^T$).*

Proof. $Z^T = (M \circ \Pi)^T = M^T \circ \Pi^T = M \circ \Pi^T$, so symmetry of \mathcal{Z} follows from that of \mathcal{P} . In addition we observe that $E(Z) = E(Z^T)$, since $R(Z) = R(Z^T)$ and

$$D(Z||Z^T) = \sum_{i>j} z_{ij} \log\left(\frac{z_{ij}}{z_{ji}}\right) + z_{ji} \log\left(\frac{z_{ji}}{z_{ij}}\right) = D(Z^T||Z).$$

Then if Z is an optimal allocation, so is Z^T and by convexity the symmetric matrix $\frac{1}{2}(Z + Z^T)$ must also be optimal. \square

B. A Markov chain approach to optimizing efficiency

The optimization of the energy $E(Z)$ over \mathcal{Z} is a convex program, but faces the difficulty that the domain description is not easily decentralized; moreover, even a ‘‘central planner’’ with global information would have to deal with large number of vertices in the polytope \mathcal{Z} . For this reason we explore here a stochastic optimization alternative (termed ‘‘Gibbs sampler’’ [3]): construct a Markov chain whose stationary distribution concentrates around configurations of minimum energy. Interestingly, this approach has been shown to yield decentralized strategies for interference-constrained wireless networks [9].

We pursue this idea in our context, first looking at efficiency alone. Define a continuous time Markov chain in the space of configurations \mathcal{X} , with transition rates

$$q(X, X + e_{ij}) = W_0 \exp\left(\frac{\mu_{ij}}{T}\right) \mathbb{1}_{\{X+e_{ij} \in \mathcal{X}\}} \quad (16)$$

$$q(X, X - e_{ij}) = \mathbb{1}_{\{X-e_{ij} \in \mathcal{X}\}} \quad (17)$$

Here e_{ij} denotes the matrix with a single ‘1’ in entry (i, j) , so transitions only add one new link or turn off an existing one. The parameter W_0 reflects the aggressiveness to occupy the medium, associated in practice with the length of idle time slots. μ_{ij} is as usual the link capacity, so we are exponentially favoring turning on faster links; this is moderated by the global ‘‘temperature’’ parameter T . Transitions that turn off a link have common rate equal to unity, which sets the global time scale. We take $W_0 = 1$ below, see [21] for the general case.

Proposition 6. *The Markov chain defined by (16)-(17) is time reversible and has invariant distribution (for $W_0 = 1$)*

$$\pi_T(X) = \frac{\exp\left(\frac{1}{T} \sum_{ij} x_{ij} \mu_{ij}\right)}{C_T}, \quad (18)$$

where C_T is a normalizing constant.

Proof. It suffices to show that the given π_T verifies the detailed balance equations

$$\pi_T(X)q(X, X + e_{ij}) = \pi_T(X + e_{ij})q(X + e_{ij}, X)$$

or equivalently

$$\frac{\pi_T(X + e_{ij})}{\pi_T(X)} = \exp\left(\frac{\mu_{ij}}{T}\right) = \frac{q(X, X + e_{ij})}{q(X + e_{ij}, X)}.$$

\square

Noting that $\sum_{ij} x_{ij} \mu_{ij} = R(X \circ M)$ is the total throughput of configuration X , we see that the probability distribution in (18) is concentrated on the most efficient configurations. This effect becomes more dramatic as $T \rightarrow 0$ (at the expense of a longer time to reach steady state).

C. Reciprocity in the Markov approach

As argued before, however, in a P2P setting we are not satisfied with efficiency, and seek some measure of reciprocity as well. Since our energy $E(Z)$ in (12) reflects this, we could aim for a steady-state distribution where the exponent $R(X \circ M)$ in (18) is replaced by $-E(X \circ M)$.

A first difficulty arises with our reciprocity measure: since it is common (e.g. in the WiFi case discussed) for links (i, j) and (j, i) to interfere, the KL divergence would be $D(Z||Z^T) = \infty$ for $Z = X \circ M$, rendering our energy useless. KL divergence is well adapted to motion in the *interior* of the feasible set, as was highlighted in Section IV; but not for random motion along its boundary as proposed here.

In our earlier work on the Gibbs sampler for wired P2P networks [19], [20], an alternative quadratic measure was used to impose peerwise reciprocity:

$$\|Z - Z^T\|_F^2 = \sum_{ij} (z_{ij} - z_{ji})^2. \quad (19)$$

But even if this latter quadratic measure does not blow up, it is still of limited use in this situation. Consider again the case where links (i, j) and (j, i) interfere, then $Z = X \circ M$ and its transpose Z^T are *orthogonal* matrices with the Frobenius inner product inherent in (19), for any $X \in \mathcal{X}$. So

$$\|Z - Z^T\|_F^2 = \|Z\|_F^2 + \|Z^T\|_F^2 = 2 \sum_{ij} \mu_{ij}^2 x_{ij}.$$

This quantity rewards efficiency (in a quadratic way) rather than any form of reciprocity.

The underlying difficulty is that peerwise reciprocity in this case cannot be measured by the *current* configuration, since necessarily there will be mutual imbalance at any given time. We need more memory in the system, keeping track of ‘‘past performance’’ as a way of guiding our reciprocity dynamics.

An approach in this direction was proposed in the thesis [21]. Peers maintain an aggregate discrepancy measure

$$d_{ij}(t) = \int_0^t [z_{ij}(\tau) - z_{ji}(\tau)] d\tau, \quad (20)$$

and use it to modulate the transitions of the Markov chain.

In particular, the transition rate in (16) is modified to

$$q(X, X + e_{ij}) = W_0 \exp\left(\frac{\mu_{ij}}{T} - \beta d_{ij}(t)\right) \mathbb{1}_{\{X+e_{ij} \in \mathcal{X}\}}, \quad (21)$$

where $\beta > 0$. This discourages opening connections to peers who “owe” reciprocity bandwidth; furthermore, the integral action present in (20) means the $d_{ij}(t)$ will not stabilize unless there is symmetry between Z and Z^T in a mean sense.

While analyzing such time-varying stochastic dynamics seems hard, some conclusions are obtained in [21] by assuming a separation of time-scales, where d_{ij} varies slowly with respect to the Markov dynamics. In particular:

- For fixed d_{ij} , the Markov chain is assumed to reach its stationary distribution, which for $W_0 = 1$ takes the form

$$\pi(X) = \frac{\exp\left(\sum_{ij} x_{ij} \left(\frac{\mu_{ij}}{T} - \alpha d_{ij}\right)\right)}{C_T}.$$

- The expected rates $\mathbb{E}(Z)$ are computed in steady-state, and used to drive a deterministic dynamics of the form (20). An equilibrium with strong reciprocity $Z^* = (Z^*)^T$ and rewarding efficiency is found. Its global stability is established through Lyapunov methods.

The thesis [21] also contains Matlab simulations, exhibiting the performance of this method in terms of reciprocity and efficiency. The effect on temperature on connection diversity is also investigated. Some limitations are noted:

- (i) One concern is the long convergence time (thousands of iterations reported in [21]), consistent with the slow-scale analysis available. For this approach to be practical one must imagine that the Markov dynamics *replaces* the standard 802.11 medium access, and thus works at the millisecond scale, whereas reciprocity factors d_{ij} are updated at the second scale as is standard practice in P2P.
- (ii) The method enforces strong peerwise reciprocity over time, with efficiency a secondary objective. If transmission rates are symmetric this is justified by Proposition 5, otherwise a more gradual tradeoff seems more adequate.

In this regard, it appears one could overcome both limitations with transition rates of the form (21), but where the discrepancy factors d_{ij} are computed based on the last few (say L) exchanges. Formally, this means the Markov state is given by the recent configurations $(X(t), X(t-1), \dots, X(t-L))$; this unfortunately does not yield a reversible Markov chain. So while it may behave well in practice, with the parameter L serving as a knob for a more gradual efficiency-reciprocity tradeoff, analytical studies do not appear straightforward.

VI. CONCLUSION

In this paper we have investigated the dual objectives of efficiency and reciprocity in P2P networks, successively incorporating features of a wireless substrate: multiple-rate physical layers and interference. In the first case, we have developed a decentralized reciprocity mechanism that optimizes a tradeoff between overall throughput and peerwise KL

divergence; while some details of the proof are pending, we have strong evidence of a fast convergence to optimality. In the interference case, decentralization was pursued through a stochastic algorithm that regulates multiple access control (CSMA, under CTS/RTS) by rewarding efficiency and penalizing imbalance when opening connections. We presented a summary of analytical results are available in [21], and laid out a suggestion for future investigations.

ACKNOWLEDGMENT

The authors were partially supported by AFOSR-US under Grant FA9550-12-1-0398.

REFERENCES

- [1] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation*. Englewood Cliffs, NJ: Prentice Hall, 1989.
- [2] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, UK: Cambridge University Press, 2004.
- [3] P. Brémaud, *Markov Chains: Gibbs Fields, Monte Carlo simulation and queues*. New York, NY: Springer, 1999.
- [4] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, “Layering as optimization decomposition: A mathematical theory of network architectures,” in *Proc. of the IEEE*, vol. 95, Jan 2007, pp. 255–312.
- [5] B. Cohen, “Incentives build robustness in BitTorrent,” in *Proc. 1st Workshop on the Economics of Peer-to-Peer Systems*, 2003, pp. 1–5.
- [6] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY: Wiley-Interscience, 1991.
- [7] B. Fan, J. C. Lui, and D.-M. Chiu, “The Design Trade-Offs of BitTorrent-Like File Sharing Protocols,” *IEEE/ACM Trans. on Networking*, vol. 17, no. 2, pp. 365–376, 2009.
- [8] A. Ferragut and F. Paganini, “Content dynamics in P2P networks from queueing and fluid perspectives,” in *Proc. of 24th International Teletraffic Congress, Krakow, Poland*, 2012.
- [9] L. Jiang and J. Walrand, “A Distributed CSMA Algorithm for Throughput and Utility Maximization in Wireless Networks,” *IEEE/ACM Transactions on Networking*, vol. 18, pp. 960–972, Jun. 2010.
- [10] F. Kelly, A. Maulloo, and D. Tan, “Rate control in communication networks: shadow prices, proportional fairness and stability,” *Journal of the Operational Research Society*, vol. 39, pp. 237–252, 1998.
- [11] D. Levin, K. LaCurts, N. Spring, and B. Bhattacharjee, “BitTorrent is an Auction: Analyzing and Improving BitTorrent’s Incentives,” *ACM SIGCOMM Computer Comm. Review*, vol. 38, pp. 243–254, 2008.
- [12] F. Paganini and A. Ferragut, “PDE models for population and residual work applied to peer-to-peer networks,” in *Proc. of 46th Annual Conference on Information Sciences and Systems (2012)*, 2012, pp. 1–6.
- [13] D. Qiu and R. Srikant, “Modeling and performance analysis of BitTorrent-like peer-to-peer networks,” *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 4, pp. 367–378, 2004.
- [14] R. Sinkhorn, “Diagonal equivalence to matrices with prescribed row and column sums. II,” *Proceedings of the American Mathematical Society*, vol. 45, pp. 195–198, Aug. 1974.
- [15] M. Vojnovic and A. Proutiere, “Hop limited flooding over dynamic networks,” in *Proc. of the IEEE INFOCOM*, 2011.
- [16] F. Wu and L. Zhang, “Proportional response dynamics leads to market equilibrium,” in *Proc. of the 39th ACM Symposium on Theory of Computing (2007)*, 2007, pp. 354–363.
- [17] X. Yang and G. De Veciana, “Fairness, incentives and performance in peer-to-peer networks,” *Performance Evaluation*, vol. 63, no. 3, pp. 175–194, 2006.
- [18] J. Zhu and B. Hajek, “Stability of a peer-to-peer communication system,” *IEEE Trans. on Information Theory*, vol. 58, no. 7, pp. 4693–4713, 2012.
- [19] M. Zubeldía, A. Ferragut, and F. Paganini, “Proportional fairness in heterogeneous peer-to-peer networks through reciprocity and Gibbs sampling,” in *Proc. of the 51st Allerton Conference*, 2013, pp. 123–130.
- [20] —, “Neighbor selection for proportional fairness in P2P networks,” *Computer Networks*, 2015, under review.
- [21] M. Zubeldía, “From resource allocation to neighbor selection in peer-to-peer networks,” Master’s thesis, Universidad ORT Uruguay, 2014.